Representing Contested Constitutional Questions Accurately

Drew Griffiths

May 2026

A descriptive account of how the contested legal, scholarly, practical, and political questions surrounding constitutional change should be represented accurately, keeping each domain distinct rather than allowing a finding in one to settle another.

Section 1. What this document is, and the method it uses

Questions about whether and how a nation within an existing state might become a state of its own are among the most contested in constitutional and international law. They are also among the most consistently misdescribed, and they are misdescribed by careful people, not only careless ones. The purpose of this document is not to answer the political question of whether any particular constitutional outcome should occur. It is to set out, as accurately as the sources allow, what is known, what is genuinely contested, and what is uncertain, and to do so in a way that keeps those three things distinct.

The reason such questions are so often misread is, on inspection, structural rather than partisan. People reach for a single answer to what is really a set of separate questions, and then allow a finding about one of those questions to settle the others. The discipline of this document is the refusal to do that. It rests on one observation, from which almost everything else follows: the questions involved belong to different domains, and a finding in one domain does not automatically carry across into another.

This is worth stating carefully, because it is the document's organising idea. The legal question of what is prohibited, the legal question of what is positively authorised as a right, the practical question of what would actually succeed, and the political question of what ought to be done are not four points on a single scale running from "impossible" to "inevitable." They are closer to different axes, orthogonal to one another, in the sense that a finding along one does not project onto the others at all. A position on one of them does not fix a position on the others, and treating them as though it does is the single most common error in this area. Legal openness is not practical availability. Practical constraint is not legal prohibition. Democratic legitimacy is not constitutional competence. Each of these is a real and separate thing, and the relationships between them have to be argued, not assumed.

From this one observation, three specific inferential errors follow, and because they are the errors this document is built to avoid, it is better to name them at the outset than to leave the reader to infer the method from the prose. They are stated here as things the document will not do, so that every later section can be checked against them.

The first is the inference from *not prohibited* to *viable*. That international law does not forbid a thing does not establish that the thing is practically available, or likely to succeed. Absence of a prohibition is a fact about the legal domain; viability is a question about the practical domain; and the first does not deliver the second.

The second is the inference from *contested* to *equally probable*. That a question is genuinely disputed, in law or among scholars, does not mean the possible answers are equally likely, or that

the practical outcome is evenly balanced. A matter can be authentically contested at the level of doctrine while the practical result remains heavily weighted in one direction. Contestation is real; it is not the same as an even distribution of probability.

The third is the inference from *authoritative now* to *settled for all time*. That a court has ruled, or that a pattern has held for decades, is strong evidence about how the relevant order presently operates. It is not evidence that the order could never change. Constitutional and international arrangements evolve: legislatures alter settlements, courts reinterpret, recognition practices shift. Treating a present authority as a permanent closure is itself an error, and, importantly, it is an error that would tend to run in only one direction, hardening whatever currently constrains. The document therefore treats judicial and historical authority as what it is: authoritative about present conditions, and contingent over time.

These three are versions of a single mistake. Each takes something true about one domain, or one moment, and treats it as a conclusion about another. Naming them in advance is what allows the rest of the document to be tested against its own stated method, rather than asking the reader to trust the author's intentions.

A word on register, since it bears on how the document should be read. The claims made here are intended to sit at the level the evidence supports, neither inflated nor hedged into vagueness. Where something is known, it is stated plainly. Where something is contested, the contestation is described rather than resolved. Where something is uncertain, the uncertainty is left standing. This is not caution for its own sake. It is the view that the accurate level of claim is itself the point, and that overstatement and understatement are both failures of description.

Section 2. The four domains

The four domains named above can be set out a little more fully, because the rest of the document is organised around them and around the boundaries between them.

The first is the question of legal prohibition: what, if anything, international law forbids. This is the narrowest of the four, and as the legal section shows, the answer is more confined than ordinary argument assumes. Few things here are flatly prohibited.

The second is the question of legal authorisation: what, if anything, international law positively confers as a right, an entitlement that others are bound to respect. This is a different question from the first. A thing may be neither prohibited nor authorised; the space between those two is not empty, and a great deal turns on recognising that it exists.

The third is the question of practical outcome: what would actually happen, and what would actually succeed, given how states, courts, and international institutions behave. This domain is governed not by legal entitlement but by effectiveness, by the conduct of the existing state, and by the recognition and acceptance of others. It is the domain in which most of the operative constraint lies, and its constraints do not derive from the legal domain; they have their own sources, which the practical section sets out.

The fourth is the question of political choice: what ought to be done, what is desirable, what a community should decide. This document does not enter that domain. It is named here only to keep it separate, because political conviction is the thing most likely to leak across the boundaries and colour an account of the other three. Keeping it out is part of the method, not an evasion of the subject.

The relationship between these domains is the heart of the matter, and it is easy to misstate. They are not arranged so that openness in one produces openness in the next. The legal domain can be relatively open while the practical domain is considerably more constrained, and these two facts are not in tension, because they are answers to different questions. A common picture

imagines a single spectrum, with legal prohibition at one end and practical inevitability at the other, along which a case slides. That picture is the source of most of the confusion. The domains are better understood as distinct, each with its own evidence and its own logic, and the analytical task is to describe each accurately and to be honest about what does, and does not, transfer between them.

This is why the document is built as it is. It treats the legal questions on their own terms, the practical questions on their own terms, and the relationship between them as something to be stated explicitly rather than assumed. The result is not a verdict. It is a map of which questions have reasonably settled answers, which are genuinely contested, and which remain open, with the boundaries between them kept visible throughout.

Section 2A. A note on nation and state

One conceptual distinction needs to be set down before the analysis begins, because confusion about it can quietly distort everything that follows. It is the distinction between a nation and a state.

A nation is a community understood to share an identity, history, culture, and sense of collective political existence. A state is a legal and institutional entity with the attributes of statehood in international law. The two are not the same thing, and they do not stand or fall together. A nation can exist without being a state; states can contain more than one nation; the existence of the one is not the legal existence of the other.

This document takes the existence of nationhood as a given where it applies, and does not treat it as a question to be argued. Nothing here turns on querying whether a community is a nation. The point of the distinction is narrower and purely analytical: the questions this document examines, what is prohibited, what is authorised, what would succeed in practice, are questions about routes to *statehood*, and they are governed by the legal and practical considerations set out above. They are not answered by nationhood. A community's status as a nation, including any history of prior sovereignty, is relevant to its identity and to the political domain, and may matter a great deal there. It does not, by itself, determine the legal or practical route to statehood under present conditions, which is the subject of the sections that follow.

Keeping nation and state distinct is therefore not a comment on either. It is simply another instance of the document's general method: two things that are often run together are in fact separate, and the analysis is clearer when they are kept apart.

Section 3. Contested Legal Questions: What International Law Prohibits, Authorises, and Leaves Unresolved

The legal domain is the one most often misdescribed in public argument, and it is misdescribed in both directions at once. Some accounts treat the law as a closed door, as though international law settles that a nation within an existing state simply cannot become a state of its own. Others treat it as an open invitation, as though the absence of a clear prohibition amounts to a recognised right. Neither is accurate, and the inaccuracy in each case comes from the same source: collapsing three distinct legal questions into one.

Those three questions are separable, and keeping them apart is the whole of the task in this section. International law may prohibit a thing, may positively authorise it as a matter of right, or may do neither and simply leave it unresolved, to be settled by facts and conduct rather than by legal entitlement. These are not three points on a single scale running from "forbidden" to "permitted." They are three different relationships between law and an event, and an accurate account has to say which relationship applies to which question. What follows takes each in

turn.

3a. What international law prohibits

The first question is the narrowest, and the answer is more confined than political argument usually assumes. International law contains strikingly little that prohibits a nation from declaring independence.

The clearest authority on this point is the International Court of Justice's 2010 advisory opinion on Kosovo. The Court was asked whether Kosovo's declaration of independence was in accordance with international law, and its central finding was a narrow one: the declaration did not violate international law, because international law contains no general prohibition on declarations of independence. The reasoning matters as much as the conclusion. The Court did not find that there was a right to make the declaration; it found that there was no rule against making it. The act fell outside what international law forbids.

It is essential to read that holding at its actual width, and no wider. The Court did not rule that Kosovo had a right to become a state. It did not rule that other states were obliged to recognise it. It did not even decide whether Kosovo was, in fact, a state. It addressed the single question put to it, the lawfulness of the declaration itself, and many states and commentators subsequently treated the opinion as confined to Kosovo's particular circumstances rather than as a general principle. The opinion is therefore authority for a precise and limited proposition: a declaration of independence is not, in and of itself, contrary to international law.

Stating this corrects a real over-closure in ordinary debate, the assumption that such a thing "cannot happen" or is simply unlawful. As a matter of the bare prohibition, that assumption is mistaken. But the correction must not be allowed to overshoot, because the absence of a prohibition is the beginning of the analysis, not the end of it. That a thing is not forbidden tells us nothing yet about whether it is authorised as a right, or whether it would succeed in practice. Those are the next two questions, and they are different questions.

3b. What international law positively authorises as a right

The second question asks something stronger than whether an act is permitted in the sense of not being banned. It asks whether international law confers a positive right to the act, an entitlement that others are bound to respect. Here the law is considerably more defined, and considerably narrower.

The relevant body of law is the right of self-determination. International law recognises a right of external self-determination, the right of a people to determine their international status up to and including independent statehood, but it recognises that right in specific and limited circumstances: principally the colonial context, and situations of foreign occupation or alien subjugation. This is sometimes described, in shorthand, as the "salt-water" conception of self-determination, because the recognised cases have characteristically involved overseas colonial territories separated from the governing power.

Outside those circumstances, the position in the scholarship, of which James Crawford's work on the creation of states is a leading statement, is that international law does not recognise a right of unilateral secession for a sub-division of an existing state. The fact that a population within a state forms a distinct community, and that a majority of that population votes for independence, does not by itself create a legal right to secede that the parent state and the wider community of states are bound to accept. For a people within an existing, functioning democratic state, self-determination is understood as a right exercised internally, through participation in the political and constitutional life of that state, rather than as a right to external separation. This internal conception is treated as consistent with the territorial integrity of the existing state,

rather than as a standing entitlement to break from it.

The point to hold precisely is the distinction between the first box and this one. That a declaration is not prohibited (3a) does not place it in this category. The two are routinely conflated, with "international law does not forbid this" silently read as "international law authorises this as a right." It does not follow. A thing can sit outside the category of the prohibited without entering the category of the positively authorised. Most of the genuinely contested terrain sits in neither, which is the subject of the third box.

3c. What international law leaves unresolved

The third question is where the greatest part of the contested terrain lies, and it is the one least often named as a distinct category at all. Between what international law prohibits and what it authorises as a right, there is a large space that the law neither forbids nor entitles, a space it leaves unresolved.

In that space, outcomes are not determined by legal entitlement, because there is no entitlement to invoke, and they are not foreclosed by prohibition, because there is no prohibition to apply. They are determined instead by other things: by effectiveness, meaning whether a claimant actually establishes and sustains the apparatus of a functioning state over a territory, and by the conduct of the parent state, meaning whether it eventually consents to, or acquiesces in, the separation. Where neither a right nor a prohibition governs, these factual and political conditions do.

This is the hinge of the whole legal picture, and it is where the Kosovo opinion actually sits. The opinion belongs here, in the unresolved space, not in the box of positive authorisation. To say that a declaration is not prohibited is to say that the law does not close the question; it is precisely not to say that the law opens it as a right. What the law does, in this large middle territory, is supply no rule that settles the question, which is then settled instead by whether a claimant can make statehood effective and whether the parent state in the end accepts it. The legal analysis runs out here, and a different layer, the practical layer, takes over. That practical layer is the subject of a later section, and it is governed by considerations quite different from, and not as open as, the legal question examined here.

The distinction this section exists to preserve

The single most common error in this area is the movement between these three boxes as though they were interchangeable. The most frequent version treats the contents of the first box as though they belonged in the second: because a declaration is not prohibited, it is taken to be authorised as a right. A second version treats the third box as though it were the second: because the law leaves a question unresolved, the unresolved space is read as though it conferred an entitlement. Both movements manufacture a legal right where the law in fact supplies none, simply by sliding from one category to an adjacent one.

Holding the three apart is what prevents that. International law prohibits very little here; it authorises a right only in narrow and specific circumstances that do not extend to a sub-division of an existing democratic state; and it leaves the larger part of the question unresolved, to be governed by effectiveness and consent rather than by law. Each of those is a different statement, and an accurate account of the legal position is one that keeps them distinct rather than letting them collapse into a single verdict of "permitted" or "forbidden." The legal domain, properly described, is neither a closed door nor an open invitation. It is a narrow prohibition, a narrow authorisation, and a wide space where the law defers to facts it does not itself control.

Section 4. The Distribution of Scholarly Opinion

The previous section set out what international law prohibits, authorises, and leaves unresolved. This section turns to a different question: not what the law is, but how scholars disagree about it. The two are easily run together, and keeping them apart matters, because a description of where expert opinion divides is not the same as a description of where the law is open. Scholarly disagreement is real, and this section aims to represent it honestly, neither smoothing it away nor inflating it into more than it is.

To do that, one distinction has to be held throughout, and it is the organising distinction of the section. There are two quite different things a scholar can disagree about. The first is what the law permits: whether a particular doctrine exists, how far it extends, what a given authority established. Call this doctrinal disagreement. The second is what would actually happen: whether, given how states and institutions behave, a particular route would succeed in practice. Call this predictive disagreement. These are different axes, and a position on one does not fix a position on the other. A scholar can hold that a doctrine arguably exists in law while expecting that it would almost never operate in practice, and the two views are entirely consistent. Failing to separate them is the most common way that a minority doctrinal argument gets quietly inflated into an implied practical possibility, and the separation is what this section exists to maintain.

A word on weight, before the substance, because it is the hardest thing to get right here. To say that a question is contested is not to say that the contending positions are equally supported. Genuine disagreement can be lopsided: a leading view and a minority view, a settled interpretation and a challenge to it, a doctrine widely accepted and a doctrine whose very existence is disputed. An honest map of disagreement has to record not only that positions differ but how the support is distributed between them. The temptation in either direction is real. One can flatten a lopsided disagreement into a false balance, presenting a minority position as though it carried equal weight; or one can dismiss a minority position as though it did not exist at all. Both are distortions. The aim here is to state, for each contested question, both that the disagreement is genuine and how the weight actually lies.

The doctrinal axis: what scholars think the law permits

On the doctrinal questions, the pattern is consistent across the contested areas, and it is worth stating as a pattern rather than case by case. In the leading scholarship on the creation of states, of which James Crawford's is the standard statement, the recurring shape is the same: in each instance there is a majority or leading position and a minority position, and the disagreement is genuine; but the minority position is, in each instance, a minority, and in some instances it is contested at the level of whether it exists in law at all.

Consider the question of admission to international institutions, which is the cleanest example because the disagreement is so sharply defined. There is a textual argument that a state might be admitted to the United Nations by the General Assembly without a favourable recommendation from the Security Council, on the reading that the bare word "recommendation" in Article 4 need not mean a favourable one. The argument genuinely exists; it has been advanced on the basis of the Charter's drafting history. But it is a minority argument, and more than that, it is a minority argument that runs directly against the leading authority: the International Court of Justice considered precisely this question in its 1950 advisory opinion on the Competence of the General Assembly for the Admission of a State to the United Nations, and rejected it, holding that admission requires both a recommendation of the Security Council and a decision of the General Assembly. So the doctrinal disagreement here is real but lopsided in a particular way: a minority textual reading exists, and the principal authority squarely contradicts it. To

report only that "scholars disagree" would be accurate but misleading; the accurate report is that a minority reading exists against a controlling authority that rejected it.

The doctrine of remedial secession is the sharpest instance, and it requires the most care, because it is contested at a deeper level than the others. Remedial secession is the proposed doctrine that a people subjected to grave and systematic oppression by their state may, as a last resort, acquire a right to secede that they would not otherwise have. Here the disagreement is not merely about the doctrine's scope but about whether it exists in international law at all, a debate that surfaced sharply in the submissions of states and in the separate opinions in the International Court of Justice's Kosovo proceedings. Some scholars treat it as an emerging or arguable principle; others regard it as unsupported, with Katherine Del Mar's treatment going so far as to call it a legal myth. The disagreement is in effect over whether remedial secession is law as it stands or merely law as some argue it ought to be. This is a genuine doctrinal debate, and it should be represented as one. But representing it honestly means stating two things at once: that the doctrine has serious proponents, and that its legal existence is disputed at the foundational level rather than settled and merely bounded. It also means stating the threshold accurately. Where the doctrine is entertained at all, the conditions attached to it are extreme: gross, systematic, ongoing violations of the most serious kind, with secession available only as a last resort where no internal remedy exists. That threshold is a feature of the doctrine as its proponents frame it, and stating it is description, not application; it is not a comment on any particular case.

The pattern, then, on the doctrinal axis is not that the contested questions are evenly balanced. It is that in each, a genuine debate exists, with a leading position and a minority one, and that the minority positions range from "arguable but outweighed" to "disputed at the level of its very existence." Recording the debate without recording that distribution would be the false-balance error; denying the debate exists would be the dismissal error. Both are avoided by stating the distribution plainly.

The predictive axis: what scholars think would actually happen

The predictive axis is different in kind, and on it the disagreement is, if anything, narrower than on the doctrinal axis. This is itself a significant fact, and it is the point at which the two axes must not be confused.

A scholar may regard a doctrinal argument as genuinely arguable, and may even find it persuasive as a matter of legal theory, while expecting that, given how states and international institutions actually behave, it would be most unlikely to operate. The admission example illustrates this exactly. Even a scholar attracted to the textual argument about Assembly admission would, as a predictive matter, have to reckon with the uniform practice that has followed the 1950 opinion: no state has been admitted in the absence of a Council recommendation, and the incentives of the permanent members to preserve the Council's gatekeeping role are strong and stable. So the doctrinal disagreement, such as it is, does not carry across into a corresponding predictive disagreement. On the predictive question, the expectation is largely uniform, whatever one's view of the textual argument.

The same divergence appears with remedial secession. Even among scholars sympathetic to the doctrine as a matter of principle, there is little expectation that it operates as an effective legal entitlement in practice. The practice that might be invoked is thin, and where the doctrine is thought to have any effect, that effect is generally said to run through recognition by other states rather than through the operation of a legal right. So once again the doctrinal disagreement, which is genuine, is not matched by a predictive disagreement of the same size. A scholar can hold that the doctrine arguably exists and simultaneously expect that it would seldom or never deliver a functioning state in practice. The two positions sit together without contradiction.

This is the heart of why the two axes must be kept apart. The doctrinal disagreements, where they exist, are real, and some of them are substantial. But they do not propagate automatically into predictive disagreements of equal size, because the predictive question is answered by a different body of evidence: the observed behaviour of states and institutions, the incentives that shape recognition and membership, the record of how comparable situations have actually resolved. On that evidence, expert expectation tends to converge more than the doctrinal debate alone would suggest. The existence of a lively argument about what the law might permit is not evidence of a correspondingly lively argument about what would actually happen.

What this section establishes

The distribution of scholarly opinion, then, has a shape, and the shape is the point. On the doctrinal axis, there are genuine disagreements, unevenly weighted, ranging from minority readings that run against controlling authority to a doctrine whose legal existence is itself disputed. On the predictive axis, expert expectation is narrower, because it is disciplined by the observed behaviour of states and institutions rather than by the openness of legal argument. And crucially, the two axes are not the same: a genuine doctrinal disagreement can coexist with a narrow predictive consensus, and frequently does.

To represent this honestly is to refuse two opposite temptations. It is to refuse to flatten the doctrinal disagreements into a false balance, as though every contested question were a coin-toss among equally weighted views. And it is equally to refuse to erase them, as though the minority positions did not exist or were unworthy of statement. The disagreements are real; they are unevenly weighted; and they live mostly on the doctrinal axis rather than the predictive one. Stating all three of those things at once is what an accurate account of the scholarship requires, and it is what the later treatment of practical constraints, which turns on the predictive evidence rather than the doctrinal debate, then builds upon.

Section 5. Practical Constraints and Recognition Realities

The previous section ended by handing the question over. Where international law neither prohibits nor authorises, it leaves the outcome to be settled by effectiveness and by the conduct of the parent state. That handover now has to be made good. It is not enough to assert that the practical layer is more constrained than the legal one; that relationship has to be shown, on its own evidence, or it is merely a claim borrowed from the section before. This section sets out that evidence.

A word first on what kind of claim this section makes, because it governs everything that follows. The constraints described here are statements about how the present international and constitutional order operates. They are drawn from the record of how states, courts, and international institutions have actually behaved, and from the incentives that shape that behaviour. They are therefore strong evidence for what is operationally likely under present conditions. They are not claims that the order could never change, that the record fixes the future, or that what has held since 1945 is a law of nature. Constitutional and international orders evolve: legislatures alter settlements, recognition patterns shift, institutions reform, courts reinterpret. The claim of this section is about present operational reality, which is a different and more defensible thing than a claim about permanent impossibility. That distinction is held throughout, and stated here once so it need not be repeated at every turn.

The empirical record

Begin with the plainest piece of evidence, because it does the most work and depends on the least theory. In the scholarship on the creation of states, of which James Crawford's is a

leading statement, the recorded position is that no new state formed since 1945 outside the colonial context has been admitted to the United Nations over the opposition of the predecessor state. Across roughly eight decades, separations of a constituent part of an existing state have ultimately involved the consent, or the eventual acquiescence, of the state being left. Where a parent state has firmly opposed a separation and not relented, a durable, recognised, fully participating new state has not emerged from that opposition.

This is an empirical record, not a rule, and it should be read as exactly that. It does not prohibit anything; the previous section established that the bare act of declaring is not prohibited. What the record shows is a different thing: that the conditions under which separations have actually succeeded have characteristically included the parent state's consent or eventual acquiescence, and that unilateral separation against sustained opposition has not, in this period, produced a functioning recognised state. The record is evidence of a strong operational regularity under present conditions. It is not a metaphysical bar, and it is not offered as one.

The regularity holds across the period regardless of the merits of particular claims, which is part of what makes it evidence about the system rather than about any one case. It reflects something structural in how the present order processes these questions, and the rest of this section is in effect an account of what that structural something is.

Effectiveness, and why it is not as available as it sounds

In the space the law leaves unresolved, the criterion most often named as decisive is effectiveness: whether a claimant actually establishes and sustains the apparatus of a functioning state over its territory. At first encounter this can sound like an open door, a test that a determined claimant could simply meet by building and holding the institutions of statehood. The record suggests it is considerably less available than it sounds, for a reason that connects effectiveness to the rest of the section.

Effectiveness is not assessed in a vacuum. A claimant's ability to establish and sustain functioning institutions over a territory depends heavily on whether the parent state contests that control and on whether other states treat the claimant as what it claims to be. Where a parent state actively opposes a separation, the practical establishment of effective, uncontested control becomes correspondingly harder. And even where some degree of control is established, the capacity to function as a state, to enter into the relations and arrangements that statehood consists of in practice, depends on others' willingness to deal with the entity as a state. Effectiveness, in other words, is entangled with consent and recognition rather than independent of them. It is a real criterion, but it is not a self-contained route that bypasses the conduct of the parent state and the wider community. This is why entities that have asserted control over a territory can nonetheless remain in a long-term condition of limited or absent recognition, neither extinguished nor functioning as ordinary states.

Recognition, and the doctrinal debate that does not change the outcome

The reason effectiveness and consent matter as much as they do becomes clearer at the level of recognition, which is where the practical layer is ultimately decided and which explains why the empirical record takes the shape it does.

There is a genuine and longstanding doctrinal debate about what recognition is. On the constitutive theory, recognition by other states is itself what creates a state as a subject of international law: an entity becomes a state by being recognised. On the declaratory theory, statehood exists once certain objective criteria are met, regardless of recognition, and the act of recognition merely acknowledges a fact that already obtains. The objective criteria are conventionally those of the Montevideo Convention: a permanent population, a defined territory, a government, and

the capacity to enter into relations with other states.

It is worth noticing that each theory can be made to sound favourable to a would-be state, which is a signal that the doctrinal debate is not where the practical question is actually settled. The declaratory theory can sound favourable because it makes statehood objective: meet the criteria and you are a state, with recognition a mere formality. The constitutive theory can sound favourable from the opposite direction, because it treats recognition as genuinely creative and important. But the constitutive theory is in fact unfavourable to an unrecognised claimant, because on that theory an entity that is not recognised is not a state at all, and states retain considerable discretion over whether to recognise. The fact that both theories can be enlisted on the optimistic side, in opposite directions, is the clue that neither theory alone determines the operational reality.

What determines it is practice, and the settled scholarly view is that practice follows neither pure theory. Neither the constitutive nor the declaratory theory satisfactorily describes how recognition actually works; contemporary practice is a hybrid, shaped by legal criteria and by political and strategic considerations together. The composite position is captured in the formulation that an entity meeting the Montevideo criteria is a state, but that no state is obliged to accord it recognition. The criteria define statehood in principle; recognition remains, in practice, a decision that states take or withhold.

The declaratory theory is the majority scholarly view and the one international practice nominally supports, and it is important not to overstate the point against it. The criteria genuinely matter; they are the recognised legal benchmark, and an entity that plainly failed to meet them would have no serious claim. But the declaratory theory's apparent favourability to an unrecognised claimant largely evaporates at the operational level, and it does so through one of the criteria itself. The fourth Montevideo criterion, the capacity to enter into relations with other states, is in practice supplied by recognition. An entity that no other state will deal with as a state cannot exercise the very capacity that completes the definition. So even on the theory that treats statehood as objective, recognition turns out to be operationally decisive, because it is the mechanism through which one of the defining capacities is actually realised. This is why entities can exist for long periods having arguably satisfied the first three criteria while remaining unable to function as states: the missing element is not territory, population, or government, but the collective willingness of others to deal with them, and that willingness is recognition.

The conclusion this points to is precise, and it is the heart of the practical layer. Whichever theory of recognition one adopts, the operational reality is the same: functioning statehood depends on collective recognition behaviour and institutional acceptance, not on the unilateral satisfaction of criteria. The doctrinal debate moves; the predictive reality does not. That is why this section can treat recognition as decisive without having to settle the theoretical dispute, and it is the mechanism that explains the empirical record set out above. Separations have succeeded where they have ultimately attracted the consent and recognition that functioning statehood requires, and have not succeeded, as durable recognised states, where sustained opposition has withheld it.

Institutional membership

The same pattern appears, in a particularly clear form, in the question of membership of international institutions, which is one of the principal ways a state functions as a state in the present order. Admission to the United Nations is the clearest instance. Under the Charter, admission requires both a recommendation of the Security Council and a decision of the General Assembly. The requirement of a Security Council recommendation means that admission is subject to the permanent members' veto.

The structure here reinforces, rather than relieves, the constraints already described. There is a textual argument, raised historically, that the General Assembly might admit a state without a favourable Security Council recommendation, treating the Council's failure to recommend as something the Assembly could work around. But this argument was authoritatively rejected: the International Court of Justice advised in 1950 that admission requires both elements, and that the Assembly cannot admit a state in the absence of a Council recommendation. The route that would have bypassed the Council was closed by the Court, and uniform practice has followed the Court's view. Membership of the principal institution of the international order therefore runs through a process in which the most powerful states hold a gate, and the existence of an arguable textual reading did not translate into an operational route around that gate.

This is a useful instance precisely because it shows the difference between an argument existing and a route being available. A textual ambiguity was genuinely present; it was nonetheless resolved against the bypass by authoritative interpretation and consistent practice, under present conditions. The general lesson for this section is that the existence of an argument that something might be possible is not evidence that it is operationally available, and the two must not be run together.

State continuity, and which way it runs

A further piece of the practical picture concerns what happens to the state that is left behind, because the answer bears directly on the position of any new entity. When a state undergoes a major territorial change, a question arises whether the original state continues in existence, retaining its identity, treaties, memberships, and seat, or whether it dissolves into successor states that must each establish their own position afresh.

The doctrinal distinction between continuation and dissolution is genuinely debated. But the way it is resolved in practice is again through collective acceptance rather than through the mechanical application of a rule. The clearest modern instance is the treatment of Russia as the continuation of the Soviet Union: Russia retained the Soviet seat, including its permanent membership of the Security Council, and was accepted by other states as continuing the legal personality of the predecessor, largely by conduct rather than by any formal readmission. A comparable claim made on behalf of the former Yugoslavia was not accepted, and the entity in question had to apply for membership anew. The difference between the two outcomes lay not in a formula but in whether the wider community of states accepted the claim.

For present purposes the significance is directional. Where a large state loses a smaller part, the practical disposition of the international community has been to treat the larger remaining entity as the continuator, retaining the predecessor's identity, treaties, and memberships, while the departing part is treated as the new entity that must establish its own position, including by seeking its own recognition and its own admission to institutions. Continuity, decided by collective acceptance, has tended to favour the larger remaining state. This is not a rule that could not change, and it is not offered as one; it is the observed disposition of the present order, and it adds to rather than relieves the practical burden that falls on a new entity emerging without the parent's consent.

The domestic constitutional layer

The constraints discussed so far operate at the level of international law and the international community. A parallel pattern appears within domestic constitutional orders, and it is worth setting alongside the international picture because it shows the same distinction at work one level down.

Domestic constitutional systems characteristically distinguish between the political significance

of a democratic demand and the legal authority to act on it unilaterally. Two judicial treatments illustrate the point, and they are linked, because the later court drew on the earlier. The Supreme Court of Canada, asked about the secession of Quebec, advised that there was no unilateral right to secede under either Canadian constitutional law or international law in Quebec's circumstances, while also holding that a clear majority on a clear question would give rise to a constitutional obligation on the other parties to negotiate. The United Kingdom Supreme Court, asked in 2022 whether the Scottish Parliament could legislate for an independence referendum without the agreement of the United Kingdom Parliament, held unanimously that it could not, because such legislation related to matters reserved to the United Kingdom under the devolution settlement as it stood. In reaching that conclusion the Court addressed and rejected an argument that the domestic settlement should be read in light of an international-law right of self-determination, observing, with reference to the Canadian decision, that such self-determination claims are confined to contexts of colonial rule or foreign occupation.

The two decisions are not identical, and the difference should be stated rather than smoothed over. The Canadian court affirmatively articulated an obligation to negotiate following a clear democratic mandate. The United Kingdom court did not address any such obligation; it confined itself to the narrower competence question actually before it. So the comparison is precise rather than neat. What both orders share is the separation of two things that public argument often fuses: the political weight of a democratic demand, which both treat as significant, and the unilateral legal authority of a sub-state institution to initiate separation on its own, which neither recognises. In each order, the lawful route to a constitutional change of this magnitude runs through the central order's own processes and not around them. In the United Kingdom case this is visible not only in the judgment but in the conduct around it: the recognised lawful path to a referendum lay in an agreement between the two governments conferring the necessary competence, the mechanism used for the referendum held in 2014.

These are statements about how these constitutional orders presently allocate competence, established by reference to the settlements as they currently stand. They are strong evidence for present operational reality. They are not statements that the settlements could not be altered, that competence could not be conferred, or that the political environment could not change; indeed the existence of a lawful route through inter-governmental agreement is itself an acknowledgement that the central order can confer what it presently withholds. The point is the distinction between democratic pressure and unilateral legal competence, not a prediction that the distinction must hold forever.

What this section establishes

Drawing the threads together without overstating them. The legal layer, as the previous section showed, is relatively open: it prohibits little and resolves little, leaving a wide space to be settled by effectiveness and consent. This section has set out why that space is, under present conditions, considerably less open than its legal description alone would suggest. The empirical record of the period since 1945 shows separations succeeding, as durable recognised states, with the parent's consent or eventual acquiescence rather than against sustained opposition. Effectiveness, the criterion the law points to, is entangled with consent and recognition rather than independent of them. Recognition, whichever theory one holds about its nature, is operationally decisive and remains a discretionary collective act. Membership of international institutions runs through gates held by powerful states. State continuity, decided by collective acceptance, tends to favour the larger remaining entity. And within domestic constitutional orders, the authority to initiate separation unilaterally is characteristically withheld from sub-state institutions, the lawful route running through the central order's own processes.

These constraints are convergent: they arise at different levels, international and domestic,

doctrinal and institutional, yet they point the same way, because they share a common mechanism. The present order makes functioning statehood depend not on the unilateral satisfaction of criteria but on collective recognition and acceptance, and the conduct of the parent state weighs heavily in whether that recognition and acceptance are forthcoming. That is why the practical layer is more constrained than the legal layer: not because anything is prohibited, but because what is unprohibited still has to be made real through the acceptance of others, and that acceptance is neither automatic nor unilaterally controllable.

It bears repeating, because it is the discipline this section has tried to keep, that all of this describes the operation of the present order and the record of the recent past. It is strong evidence for what is operationally likely under present conditions, and it is offered as nothing more, and nothing less, than that.

Section 6. Why Legal Uncertainty Does Not Itself Confer Viability

The two preceding sections can now be drawn together, and the join between them is the single most important point in the document. It can be stated in two sentences. The existence of legal uncertainty does not itself create a viable route. It creates an area of contestation.

Everything in this section follows from holding those two sentences apart from the conclusions they are most often mistaken for. The legal section showed that the law prohibits little and authorises a right only narrowly, leaving a wide unresolved space. The practical section showed that this space, under present conditions, is governed by effectiveness, consent, and recognition, and that these are considerably less open than the legal description alone suggests. Section 6 does not add new evidence. It states precisely what those two findings, taken together, do and do not establish.

The firewall stated

The relationship between the legal layer and the practical layer is the relationship this document has called a firewall, and the term should be understood literally rather than dramatically. A firewall does not deny that there is fire on the other side; it prevents one thing from passing into another. Here, what must be prevented from passing is the inference that openness in the legal domain implies openness in the practical domain. Those are different domains, governed by different things, and a finding in one does not transfer to the other.

The principle can be put plainly. Legal ambiguity does not propagate outward into equal practical ambiguity. That a question is unresolved in law does not mean it is equally unresolved in practice, because the practical domain is not waiting on the legal domain to make up its mind. It is governed by its own conditions, the conditions set out in the previous section, and those conditions can be settled and constraining even where the legal position is genuinely open. The legal openness and the practical constraint coexist; the second does not dissolve the first, and the first does not soften the second.

This is why the firewall is not a closure. It does not say that the legal openness is illusory, or that the unresolved space is really prohibition in disguise. The legal openness is real, and the document has insisted on it. What the firewall says is narrower: that the openness is legal, and that legal openness is not the same thing as practical availability. The two are simply different claims about different domains.

What kind of claim this is

It matters to be exact about the nature of the practical claim, because the synthesis is the point at which a careful claim is most easily rounded up into a larger one it does not need and cannot

fully support.

The claim is operational, not metaphysical. The evidence of the previous section supports the proposition that, under present conditions, functioning statehood depends on recognition and acceptance, and not on the unilateral satisfaction of criteria alone. That is a statement about how statehood works in practice: about what an entity must secure in order to function as a state, given how the present international and constitutional order actually behaves. It is not a statement about what statehood is. The document does not claim that recognition constitutes statehood, or take a position in the theoretical dispute between the constitutive and declaratory accounts. It does not need to. The operational claim holds whichever account of statehood's nature one prefers, because even the declaratory account, which treats statehood as objective, concedes in practice that the capacity to function as a state is supplied through recognition. The narrower claim is therefore the more durable one: it survives disagreement about the deeper question, because it does not depend on resolving it.

This is also why the claim is bounded in time. It describes the operation of the present order and the record of the recent past. It is strong evidence for what is operationally likely under present conditions, and it is offered as that. It is not a claim that the order could not change, that the conditions are permanent, or that what has held is fixed. The constraints are real and presently weighty; they are not eternal. Holding the claim at its actual scope, operational and present, is not a hedge. It is the difference between a defensible predictive statement and an overreaching one, and the defensible version is the stronger of the two.

The two inferences this prevents

The firewall exists to prevent two specific movements of thought, and they are worth naming directly because they are the errors the whole document is built to resist.

The first is the inference from "not prohibited" to "viable." Because the legal section established that a declaration of independence is not, in itself, prohibited, it is tempting to read that as establishing that such a route is available. It does not. The absence of a prohibition is a fact about the legal domain; viability is a question about the practical domain; and the firewall is precisely the refusal to let the first stand in for the second. Not prohibited means not prohibited. It does not mean workable, and it does not mean likely to succeed.

The second is the inference from "contested" to "equally probable." Because several of the questions examined are genuinely contested, in law and among scholars, it is tempting to treat the contestation as if it distributed probability evenly across the possible outcomes, as if a disputed question were therefore an open one in which any answer is as likely as any other. It does not follow. A question can be genuinely contested at the level of doctrine while the practical outcome remains heavily weighted in one direction, because doctrinal contestation and predictive probability are different things. The previous sections showed this repeatedly: textual arguments that genuinely exist but have been authoritatively closed; doctrines whose very legal existence is disputed and which carry almost no practical operation; theoretical debates that move while the predictive reality does not. Contestation is real, and it is not the same as equiprobability.

Both inferences share a structure. Each takes a true statement about one domain and treats it as a conclusion about another. The firewall is the consistent refusal to make that move, in either direction, and applied to every domain alike. The third inference the document set out to resist, from present authority to permanent settlement, has the same structure across time rather than across domains, and it was addressed above in the account of what kind of claim this is: it is the refusal to read a description of the present order as a statement about all possible futures.

What the synthesis establishes, and what it leaves open

The result is a position that is easy to state and worth stating without embellishment. In the present order, the legal domain is relatively open: it prohibits little and confers a right only in narrow circumstances. The practical domain is considerably more constrained: it is governed by effectiveness, consent, and recognition, which are not automatic and not unilaterally controllable. These two facts are both true at once, and the relationship between them is not that one cancels the other but that they describe different things. Legal openness is genuine. Practical constraint is genuine. Neither is evidence about the other.

That is the whole of the firewall, and it is deliberately the whole of it. The document does not proceed from here to a verdict on whether any particular constitutional outcome should or should not occur, nor to a prediction dressed as a certainty. It establishes a way of seeing the question accurately: legal possibility, scholarly disagreement, practical likelihood, and political choice kept as the distinct things they are, with uncertainty in one never silently borrowed to colour another. Which of those questions matters most, and what follows from holding them apart, is not for this section to decide. The synthesis clarifies the structure. It does not close the matter, and it is not trying to.

Section 7. Political Strategy as a Separate Category

The four domains set out at the start included one the document named only to set aside: the political question of what ought to be done. Having worked through the legal, scholarly, and practical questions, it is worth returning to that fourth domain briefly, not to enter it, but to mark clearly where its boundary lies.

There is a real and consequential set of questions about political strategy: what a movement, a government, or a community should do, what sequence of steps might be wise or unwise, what is desirable, what is worth the cost, how competing goods should be weighed. These questions matter, and they are not lesser questions for being outside the scope of this document. They are simply different questions, governed by values, judgement, and choice rather than by the description of law and practice that this document has undertaken.

The reason for keeping them out is the same reason that has governed throughout. The political question belongs to a different domain from the legal, scholarly, and practical ones, and a finding in those domains does not settle it. That the law prohibits little, that scholarly opinion divides in particular ways, that the practical constraints run as they do under present conditions: none of these tells anyone what ought to be done. They are inputs a person might weigh in forming a political judgement, but they do not make the judgement, and this document does not make it either. To slide from the descriptive account into a recommendation would be to commit, in the other direction, exactly the kind of category transfer the whole document is built to resist.

This is not evasion, and it is not indifference. It is a recognition that the descriptive task and the political task are distinct, that doing the first well requires not pretending it settles the second, and that the credibility of the description depends on its not being bent toward a preferred conclusion. A reader who holds a strong political view, in any direction, should be able to use this account without finding that it has quietly done their deciding for them. That is only possible if the political domain is left genuinely open.

So this section names the domain and declines it. What should be done with the picture set out here belongs to political judgement, to those who hold political responsibility, and to the people they answer to. The account stops at the edge of that judgement, deliberately, and leaves it where it belongs.

Section 8. A Closing Note on Method

It remains only to state, as plainly as possible, what kind of account this has been, so that it is read for what it is.

The aim throughout has been to describe accurately, at the level the evidence supports, and to keep distinct things distinct. Where the law is open, the document has said so; where it is narrow, it has said that; where a question is genuinely contested, the contestation has been described rather than resolved or explained away; and where the practical constraints are strong, they have been set out as strong, and as what they are, features of how the present order operates rather than laws of nature. The discipline has been to let each claim sit at its own level and not to borrow conviction from one domain to settle another.

That discipline has to run in every direction at once, or it is not discipline at all. The uncertainty has been applied symmetrically. The same refusal that prevents legal openness from being read as practical viability also prevents practical constraint from being read as legal prohibition; the same caution that declines to treat a contested doctrine as settled also declines to treat it as worthless; the same care that keeps an authoritative ruling from being overread as permanent also keeps it from being dismissed as merely provisional. An account that relaxed the discipline whenever it pointed in a congenial direction would not be neutral, however balanced its conclusions happened to look. Neutrality here is not a midpoint between positions; it is the consistent application of the same standard to every claim, whichever way the claim cuts.

Two limits of the account are worth restating, because they are easy to forget once a description reads as confident. The first is that it is descriptive, not prescriptive: it sets out what is known, contested, and uncertain, and stops there, leaving what ought to be done to the political domain it has declined to enter. The second is that it is bounded in time. Everything here describes the present order and the recent record. Constitutional and international arrangements change, and nothing in this account claims otherwise; its statements are strong evidence about present conditions and are offered as exactly that, not as predictions fixed against any future.

What the document offers, then, is not a verdict but a way of seeing. It is a set of distinctions held steadily: prohibition from authorisation, legal openness from practical availability, doctrinal disagreement from predictive expectation, present authority from permanent settlement, democratic significance from unilateral competence, and description from choice. Holding those apart is the whole of the method, and the value of the account lies there rather than in any conclusion drawn from it. The structure is what is meant to remain when the reading is done.